

## Lec 21: Virtual Memory

**Kavita Bala**  
**CS 3410, Fall 2008**  
Computer Science  
Cornell University

## Announcements

---


- HW 4 out: smash the stack
  - Due Nov 14<sup>th</sup>
- PA 3 out Nov 14<sup>th</sup>
  - Due Nov 25<sup>th</sup> (feel free to turn it in early)
  - Demos and pizza party: Dec 1<sup>st</sup> or 2<sup>nd</sup>
- Prelim 2: Dec 4th
- Final project: distributed multicore ray tracer
  - Due exam week

© Kavita Bala, Computer Science, Cornell University

## How to make it work?

---

- Challenge: Virtual Memory can be slow!
- At run-time: virtual address must be translated to a physical address
- MMU (combination of hardware and software)



© Kavita Bala, Computer Science, Cornell University

## Address Translation

---


- How to translate addresses?
  - Per word? Much too expensive
  - Per block? Sure, but what is block size?
- Costs dictate granularity of translation
  - Cost to disk is very large
  - Block size has to be large too
    - Amortization
- Page Table: stores this translation
  - Basically a huge array of translations
  - Each process has one

© Kavita Bala, Computer Science, Cornell University

## Virtual Addressing with a Cache

---

- It takes an *extra* memory access to translate a VA to a PA



- This makes memory (cache) accesses very expensive (if every access was really *two* accesses)

© Kavita Bala, Computer Science, Cornell University


## Translation Lookaside Buffer (TLB)

---

- Hardware solution to problem
- A small cache of recently used address mappings
  - TLB hit: avoid a page table lookup

© Kavita Bala, Computer Science, Cornell University

## A TLB in the Memory Hierarchy



- A TLB miss:
  - If page in main memory, TLB miss can be handled (in hardware or software)
    - Load from the page table into the TLB
    - Takes 10's of cycles

© Kavita Bala, Computer Science, Cornell University


## Hardware/Software Boundary

- Translation Lookaside Buffer (TLB) that caches the recent translations
  - TLB access time is part of the cache hit time
  - May allow an extra stage in the pipeline for TLB access

© Kavita Bala, Computer Science, Cornell University

## Remove TLB from critical path?

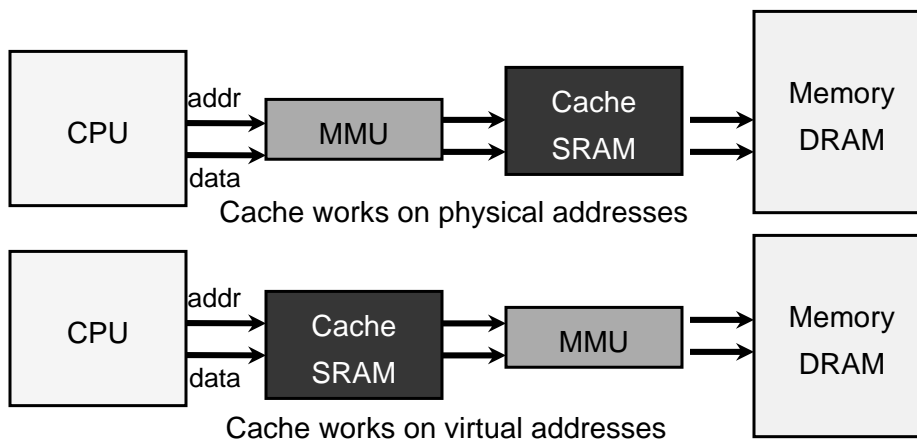
- A virtually addressed cache would only require address translation on cache misses



- Cons: have to flush the cache on context switch

© Kavita Bala, Computer Science, Cornell University


## Virtual vs. Physical Caches



- L1 (on-chip) caches are typically virtual
- L2 (off-chip) caches are typically physical

© Kavita Bala, Computer Science, Cornell University


## A TLB in the Memory Hierarchy



- A TLB miss:
  - If the page is not in main memory
    - Page fault!
    - Takes 1,000,000's of cycles to service a page fault
- TLB misses are much more frequent than true page faults

© Kavita Bala, Computer Science, Cornell University

## Paging



© Kavita Bala, Computer Science, Cornell University


## Example: paging to disk

---

- Compiler (gcc) needs a new page of memory
- Trap/exception into OS
  - OS gets page from application (vi)
- If page is clean, give up page to gcc
- If page is dirty, ... only copy in memory
  - Write to disk, before giving it to gcc
- Mark page invalid in vi
  - (if vi needs this soon, it will trap)

© Kavita Bala, Computer Science, Cornell University

## Working set model



© Kavita Bala, Computer Science, Cornell University

## Thrashing


- Thrashing: processes on system require more memory than it has
  - gcc removes vi page
  - vi comes back and removes gcc page
  - Spend lot of time waiting to read pages
  - i/o device at 100% utilization
    - But no useful work is getting done
  - We want Virtual Memory
    - Size of disk, access time of main memory
  - We have
    - Access time of disk

© Kavita Bala, Computer Science, Cornell University

## Reasons for Thrashing


- Process doesn't reuse memory

– Past != future




- Process reuses memory

– But it doesn't fit



- Individually, processes fit in memory

– But there are too many



© Kavita Bala, Computer Science, Cornell University